

A Distributed Spiking Neuron Model of Attention in the Stroop Task

Emilie Caron (emilie.caron@uwaterloo.ca)

Department of Psychology, University of Waterloo
200 University Avenue West, Waterloo, ON, Canada, N2L 3G1

Terrence C. Stewart (terrence.stewart@nrc-cnrc.gc.ca)

National Research Council of Canada, University of Waterloo Collaboration Centre
200 University Avenue West, Waterloo, ON, Canada, N2L 3G1

Abstract

We present a spiking neuron-based model of the Stroop task where the attention mechanism is entirely implemented with distributed representations. This is done by using the Neural Engineering Framework and the associated Semantic Pointer Architecture to implement a selective attention mechanism. The resulting system exhibits the Stroop effect, as well as the associated Facilitation and Interference effects. In contrast with previous models, these effects are not generated via a localist competition mechanism. Rather, these effects are a result of controlled unbinding of information from a combined distributed representation.

Keywords: Stroop task; neural engineering framework; semantic pointer architecture; spiking neurons; distributed representation

Introduction

Current models of the neural mechanisms underlying selective attention (in tasks such as the Stroop task) rely on a localist representation of concepts. That is, they postulate that there are individual and separate neurons (or groups of neurons) representing concepts such as RED and BLUE, and whether attention should be paid to the COLOR of a word or to the WORD itself. While these sorts of localist neural models are common, new methods have been developed for creating neural models that make use of distributed representations, where the color BLUE would be represented by a particular pattern of activity over a group of neurons, and a different pattern of activity would represent RED (or a concept such as COLOR or WORD, etc.). These sorts of representations match well to biology (e.g. Stewart & Eliasmith, 2012) and offer an alternative set of mechanisms for the manipulation and control of representations. Here, we apply these techniques for the first time to modelling the Stroop task.

Background

On a daily basis, individuals are tasked with allocating their attention to specific information given their situational demands. This is done by selectively choosing to focus on the relevant aspects of their situation and discarding the irrelevant ones (Bustamante, Lieder, Musslick, Shenhay, & Cohen, 2020). This behaviour is typically believed to be guided by our internal state and often explained using the Top-down Excitatory Biasing (TEB) model. Specifically, the TEB model suggests that representations of cognitive control guide task completion through heightened levels of activity in groups

of neurons associated with processing task-relevant information in relation to the levels of activity in groups of neurons associated with processing task-irrelevant information (Corbetta, Miezin, Dobmeyer, Shulman, & Petersen, 1991; Herd, Banich, & O'reilly, 2006). As a result, irrelevant information becomes less influential.

A common task in which this phenomenon is illustrated is the Stroop task (Stroop, 1935). The Stroop task is often referred to as a control task or a conflict-related task (Petersen & Posner, 2012) because individuals are required to selectively attend and respond to the pertinent information from the stimulus and ignore the impertinent information. Specifically, when completing the Stroop task, participants are instructed to respond as quickly and accurately as possible by naming the color of a word stimulus, all while ignoring the occasionally-distracting information coming from the word itself. Accordingly, the Stroop task typically consists of two kinds of stimuli: 1) congruent stimuli, where the word is a color word that matches the color in which it is presented (e.g., the color word RED presented in red) and 2) incongruent stimuli, where the color word mismatches the color in which it is presented (e.g., the color word RED presented in blue). Importantly, individuals typically take longer to name the color of the stimuli on the incongruent trials than on congruent trials. This difference in response time between congruent and incongruent word stimuli is often referred to as the Stroop effect and is commonly perceived to be a result of the automaticity of word reading which influences color naming on incongruent and congruent trials (e.g., Ashcraft, 1994; Rayner & Pollatsek, 1989). That is, when responding to congruent trials participants might be enlisting well-established reading processes to enhance the speed and accuracy of their responses, whereas on the incongruent trials participants might need to depend on cognitive control systems to reduce the reflex to read the color word and initiate the color-naming processes (Bugg, Jacoby, & Toth, 2008).

Notably, current neural models of the Stroop task (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Cohen, Dunbar, & McClelland, 1990; Cohen & Huston, 1994; Herd et al., 2006) have demonstrated this conflict-control mechanism whereby the representation of the activity in the simulated neurons differs when responding to the congruent trials versus responding to the incongruent trials. Thus, the representation of neural activity is heightened in response to congruent trials, where naming the color of the color word stim-

ulus is assumed to be supported by the automatic process of reading the word. At the same time, the representation of neural activity is weakened when responding to the incongruent trials, where naming the color of the color word stimulus is assumed to be hindered by the automatic process of reading the word (e.g., Herd et al., 2006; Bugg et al., 2008).

Critically, a common trait shared by these current models is their dependence on the notion of a localist representation. In other words, there are separate representations for each individual concept, such as RED or BLUE. These separate representations compete with each other, in that as the representation of RED becomes stronger, it reduces or inhibits the representation of BLUE (and all other competing terms). For models that make use of neural mechanisms, these separate representations often take the form of separate neurons, or separate groups of neurons, one group for each possible concept. These localist representation models are a relatively common method of examining conflict-related tasks such as the Stroop task. However, this sort of localist representation has been criticized and contrasted with distributed representations (e.g. Rumelhart, McClelland, & PDP Research Group, 1986), and these types of distributed representations can form closer connections to the underlying neurobiological implementation (e.g. Stewart, Bekolay, & Eliasmith, 2011).

With this in mind, we aim to create a novel model of the Stroop task that makes use of distributed representations. We believe this framework may provide more information regarding the function of the neural system on cognitive control. In addition to the use of distributed representations in our model, we also include a third kind of stimuli, the neutral word stimuli. The Neutral stimuli is all-too-frequently excluded from the Stroop task, despite it being useful for interpreting how much of the Stroop effect is driven by Interference (arising from the incongruent words) and how much of it is driven by Facilitation (arising from congruent words) (MacLeod, 1991). Although there is debate as to what consists of an appropriate neutral stimulus, for the sake of our paper we have opted to use non-color word stimuli which is a commonly used method in the color Stroop task.

The inclusion of the neutral stimuli allows us to break down the Stroop effect into its various parts and to explore these complexities in greater detail. Further, by implementing distributed representations within the model we are suggesting a very different mechanistic theory than other models of the Stroop Task. In particular, rather than relying on competitive inhibition operations, our model is based on computing mathematical operations on high-dimensional vectors. Since these operations can then be implemented using biologically detailed neurons, we can tie the model more closely to biological constraints. This allows for direct output of metrics such as response times, rather than using an abstract notion of time steps.

The goal of our paper is to demonstrate a neural mechanism that can produce the Stroop effect so that we might be able to evaluate in greater depth what is occurring at the neu-

ral level. By doing so, we hope to show exactly how top-down biasing can be implemented in a flexible manner. We also hope it can shed light on why the magnitude of the Stroop effect varies based on different situations, (e.g. 160-260ms in Dunbar & MacLeod, 1984 and 75-150ms in Augustinova, Parris, & Ferrand, 2019). Furthermore, we examine the relationship between the Stroop effect and its two components: the Facilitation effect (defined as the difference in performance between the congruent and neutral trials, i.e. the benefit when the word is congruent to its color) and the Interference effect (defined as the difference in performance between the incongruent and neutral trials; i.e. the penalty when the word is incongruent to its color).

Distributed Representation

Most models of language that use distributed representations make use of vectors (i.e. a list of numbers). That is, each basic term (RED, BLUE, COLOR, etc.) is a vector in some high-dimensional space. These vectors can be randomly chosen, or can be chosen to respect semantic similarity (so that the vector for RED is similar to the vector for PINK, for example). In the work presented here, all vectors are randomly chosen 512-dimensional unit-length vectors.

Importantly, these distributed representations can be combined in order to create representations of more complex structures. While there are many different mathematical frameworks suitable for forming these combinations (see Levy & Gayler, 2008 for an overview), they all generally have operators for binding and unbinding. For example, one can build a vector to represent "dogs chase cats" by computing the following:

$$S = \text{SUBJECT} \otimes \text{DOGS} + \text{VERB} \otimes \text{CHASE} + \text{OBJECT} \otimes \text{CATS}$$

This gives us a final resulting vector S which forms a distributed representation of that entire sentence. Given S , we can recover the individual parts by performing unbinding (written as binding by the inverse, noted by $^{-1}$). In this case, if we want the object of the sentence, we can compute

$$S \otimes \text{OBJECT}^{-1} \approx \text{CATS}$$

This is, of course, an approximation, and the resulting output will be less and less similar to the ideal vector for CATS as the vocabulary size and the number of terms that are combined (and the depth of the combinations) are increased.

In this work, we use circular convolution as a binding operation, and its associated pseudo-inverse (turning circular correlation into circular convolution) as the unbinding operator, as suggested in (Plate, 1995). These operations have been shown to be efficiently implementable by spiking neurons, and have been used in many neurally detailed models of cognitive behaviour, including sequential memory (Choo & Eliasmith, 2010), semantic search (Kajić, Gosmann, Stewart, Wennekers, & Eliasmith, 2017), emotional appraisal (Kajić, Schröder, Stewart, & Thagard, 2019), and spatial representation (Lu, Voelker, Komer, & Eliasmith, 2019). This approach is known as the Semantic Pointer Architecture, and has been

shown to scale efficiently to human-sized vocabularies and sentence structures (Eliasmith, 2013).

Distributed Representation of the Stroop Task

While distributed representations and Semantic Pointers have been used to model a variety of tasks, they have not yet, to our knowledge, been used to model the Stroop task. To do this, we take an approach that closely follows the binding and unbinding ideas described in the previous section.

First, we are not modelling the entire visual processing system. Instead, we assume that visual processing is inputting to our model a combined representation of the word and its color (and any other visual information could also be included). That is, the input representation is a vector that combines the vector for the word with the vector for the color. For example, the word RED written in blue would be presented to the model as:

$$\text{vision} = \text{WORD} \otimes \text{RED} + \text{COLOR} \otimes \text{BLUE}$$

This can be thought of as the output (final layer) of a standard deep-network vision system which is external to the model presented here, and in future work we will be including this vision system as part of the model.

In the normal Stroop task, participants are asked to identify the color of the stimuli they are seeing, but it is also possible to ask them to identify the words. This means there must be a way to dynamically change the aspect of the visual stimuli to which participants are paying attention. In our model, this is a separate input. If they are perfectly paying attention to the color, then the vector input to this part of the model is the vector COLOR.

This approach allows us to explore one possible method for producing the Stroop effect: attention might not be perfect. Instead, the attention value might be $\text{attention} = 0.7\text{COLOR} + 0.3\text{WORD}$, indicating that most of the attention is on the color of the word, but some attention is on the word itself.

Given this setup, the model itself just needs to compute the result of $\text{vision} \otimes \text{attention}^{-1}$. This results in a vector that should be close to the concept that determines the response.

However, in order to produce a response time out of our model, we also need to include a mechanism that will take the resulting vector and turn it into an explicit response (saying a color name, or pressing a button). To model this, we use a set of independent accumulators (Gosmann, Voelker, & Eliasmith, 2017). That is, we take the result value and feed it into a set of evidence accumulators, one for each possible response. This means there is a group of neurons whose activity increases based on similarity between the result and the vector for RED, another for BLUE, and so on. (Similarity is computed as the dot product). When this activity reaches a threshold, we consider that to be the time that a response is generated. Making a response suppresses activity in all the other accumulators, but there is no interaction between the accumulators until that decision is generated.

It should be noted that the only part of this model which uses localist representations is the independent accumulators used to make the final decision. All other parts of the model are distributed, and we do not make use of any direct competition mechanism, unlike the localist models discussed above.

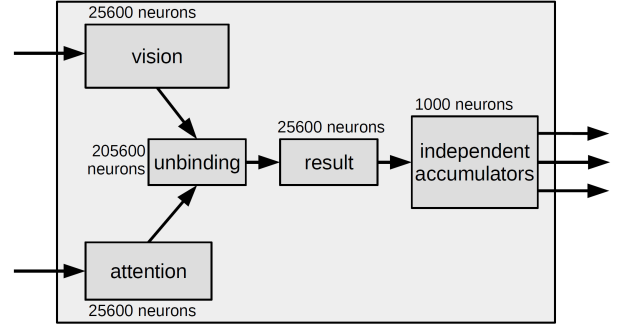


Figure 1: The model of attention in the Stroop task. Inputs are vectors, such as $\text{WORD} \otimes \text{RED} + \text{COLOR} \otimes \text{BLUE}$ for the vision and $0.7\text{COLOR} + 0.3\text{WORD}$ for the attention. Boxes represent single layers of neurons and arrows represent all-to-all connection weight matrices. Connection weights are set to optimally approximate the desired unbinding operation (circular correlation), and the independent accumulators are a set of neural integrators that build up evidence over time until reaching a threshold and producing an output response.

Neural Implementation

Previous work has shown how spiking neurons can be connected such that they represent vectors and compute functions such as circular convolution (Eliasmith, 2013), and we follow the same approach here. To have a group of neurons represent a vector, we randomly assign each neuron a “preferred direction vector”. This is the vector for which it will fire most strongly, consistent with the preferred stimuli found for many sensory and motor neurons (e.g. Georgopoulos, Schwartz, & Kettner, 1986). Each neuron is also given a random gain and bias current, providing heterogeneity in the population coding. In this situation, any vector will result in a different pattern of neural activity in the population, and we can think of this as much like a single-hidden-layer neural network where the input weights are randomly chosen.

In order to compute functions using these distributed representations, we solve for the connection weights that will lead to one group of neurons causing the desired neural activity in another group of neurons. For example, if one group of neurons represents x and another group of neurons represents y and we want $y = f(x)$, then we need to find the set of connection weights from the x population to the y population that achieves this for all x values. This can be solved using a variety of optimization techniques; here we treat it as a least-squares minimization problem and solve for the optimal connection weights. This general process is known as

the Neural Engineering Framework (NEF; Eliasmith & Anderson, 2003).

This resulting network is depicted in Fig. 1, and its behaviour is in Fig 2. The neural activity (i.e. individual neuron spikes) of a subset of the neurons in each area are shown. For the two input areas (visual and attention), we also show the computed vector that is being used to stimulate those neurons. All other neural activity is purely the result of the computed connection weights that implement the unbinding and accumulation operations. For the accumulator neurons, we also present the aggregate activity (red and blue lines) encoding the gradual increase in evidence for a decision, and the final decision that is made once a threshold is reached. The bottom row of Fig 2 shows a response time of 340ms for the incongruent case, while the congruent case is faster at 240ms.




Figure 2: Model behaviour for a congruent and an incongruent example. Grey background is spiking neural activity from randomly chosen neurons. Overlaid text is the vector presented as input. Red and Blue lines show the accumulation of the similarity between the neural activity in `result` and the ideal neural activity for `RED` and `BLUE`, respectively. The decision is made when an accumulator reaches a threshold, which then suppresses the other accumulators.

Results

In this paper, we examine the overall behaviour of our model, rather than fitting it to a particular study. The magnitude of

Stroop effects have been shown to be sensitive to a wide variety of factors (e.g. Dunbar & MacLeod, 1984; Augustinova et al., 2019). In future work we will be examining these factors and mapping them into our model, but here we just show the basic exploratory results. In general, most Stroop studies find the fastest response times for congruent trials (e.g. the word `RED` written in the color `RED`), and the slowest response times for incongruent trials (e.g. the word `BLUE` written in the color `RED`), with neutral trials somewhere in-between (e.g. the word `HOUSE` written in the color `RED`). This basic pattern is replicated by our model (Fig. 3).

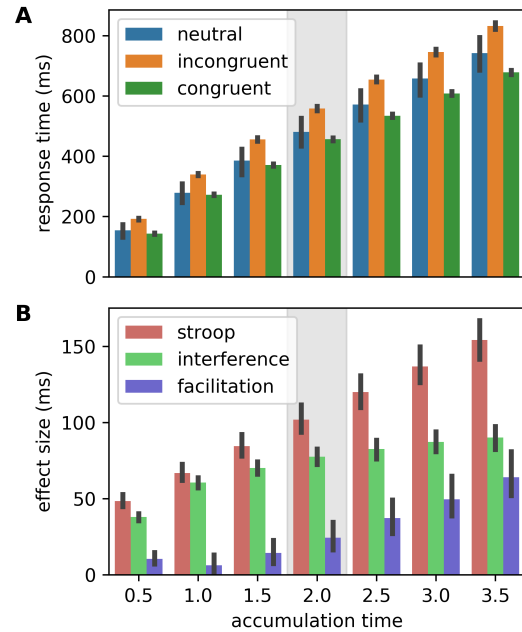


Figure 3: Model response times when varying the time over which the accumulators build up evidence. (A): Raw response times for neutral, incongruent, and congruent cases. (B): Effect sizes computed as differences in response times. Stroop = Congruent - Incongruent; Interference = Incongruent - Neutral; Facilitation = Neutral - Congruent. Shaded area shows the parameter setting used in the rest of the paper.

The actual magnitude of the response times is strongly dependent on the rate of accumulation used to make the final decision. Fig. 3A shows this linear increase in overall response time as we increase the accumulation time parameter for the independent accumulators. However, if we present the same data in a different way, Fig. 3B shows a surprising result. Here, we measure the Stroop effect (the difference in response times between congruent and incongruent trials), the Interference effect (the difference in response times between neutral and incongruent trials), and the Facilitation effect (the difference in response times between congruent and neutral trials). The Stroop effect will always be the sum of the Interference and Facilitation effects. Fig. 3B shows that increasing the accumulation time increases the Stroop effect linearly, but

the Interference effect stops increasing at around 80ms, and after that the Facilitation effect is primarily responsible for the increase in the Stroop effect. Another way to think of this is that in situations with fast response times (i.e. low amounts of accumulation time), the Stroop effect is almost entirely driven by the Interference effect, but with longer response times the Facilitation effect becomes stronger.

While we are still conducting a literature review as to how the Stroop, Interference, and Facilitation effects vary across experimental conditions, so that we can perform parameter fitting to those conditions, it should be noted that our recent experimental work has shown Interference effects in the 100-140ms range, and Facilitation effects in the 0-10ms range, for conditions where subjects are standing or sitting (Caron et al., in press).

In addition to the accumulation time, we also explored one other parameter of our basic model. The neural representation for the attention signal was initially set to $0.7\text{COLOR} + 0.3\text{WORD}$. In Fig. 4 we vary this ratio. Interestingly, while this also linearly increases the Stroop effect, it causes *no change at all* to the Interference effect. Instead, this purely changes the Facilitation effect. Furthermore, if this ratio is made to be too extreme (i.e. if the attention signal becomes $0.8\text{COLOR} + 0.2\text{WORD}$), then the model starts producing a *negative* Facilitation effect (i.e. congruent trials become slower than neutral trials). This is not a phenomenon that is commonly seen in the behavioural literature, which either indicates a lower bound on this parameter in humans (i.e. humans don't adjust their attention to be that extremely focused on color), or an indication that other features need to be added to the model.

As an attempt to add another feature to the model that might increase the Facilitation effect (and stop it from becoming negative), we also tried adding a direct connection between the visual neurons and the result neurons, bypassing the unbinding system. Biologically, these would be connections from the visual system that cannot be modulated by the attention system. This can be thought of as a sort of "automaticity", in that these connections are always feeding the WORD information to the result, no matter where high-level cognition is directing attention. The result of adding this parameter is shown in Fig 5.

With this parameter, the Interference effect is mostly unchanged, and it caused the desired increase in the Facilitation effect. However, after a certain point, it starts causing a strong drop in the Interference effect. Oddly, this does not correspond to any decrease in accuracy, and we are still analyzing the system to determine why this may be happening.

All other parameters in the model were left at the default values that have been used in previous Nengo and SPA models. Future work will analyze the effects of these other parameters, such as the number of neurons and the vector dimensionality.

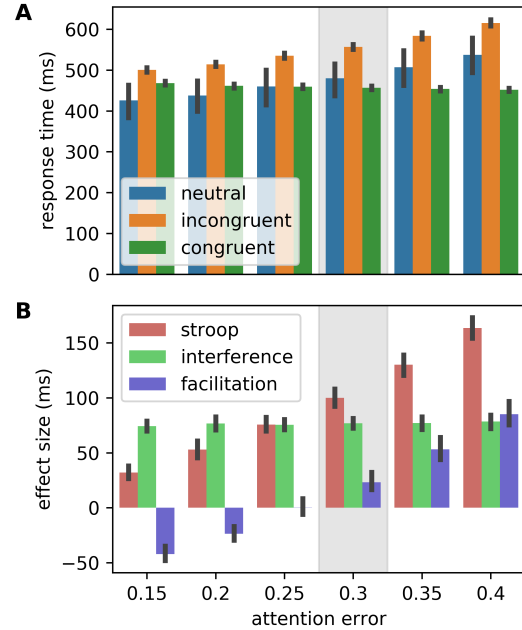


Figure 4: Model response times when varying the accuracy of the attention representation. For an attention error of x , the neural activity in the attention neurons is set to represent the vector $(1 - x)\text{COLOR} + x\text{WORD}$.

Discussion and Future Work

We have presented a neural model of the Stroop task that uses a very different mechanism than previous models. Rather than relying on a localist representation and a competition-based mechanism to exhibit the Stroop effect, we produce the Stroop effect by forming distributed representations that bind together color and word information, and implement an attention computation that extracts out the desired information. As we show, this also produces a Stroop effect. We show that this can all be implemented in spiking neurons using the same mechanisms that have been instrumental in implementing other tasks, such as sequential memory, semantic search, emotional appraisal, and spatial relations.

However, we have only begun to analyze this model and cannot yet claim that it is an improvement over previous approaches. The actual magnitude of Stroop effects varies significantly in different conditions, and we have not yet begun to map different conditions into different parameter settings for our model. Our initial parameter exploration indicates that the magnitude of the Interference effect is, surprisingly, generally unaffected by the amount of attentional error or direct automaticity, and is instead only affected by the rate of accumulation of evidence needed to make a decision, and it seems to have a soft maximum of just under 100ms.

While our primary future work is to explore the parameter space of this model and fit it to various Stroop conditions, there are also clearly many additions still needed to the

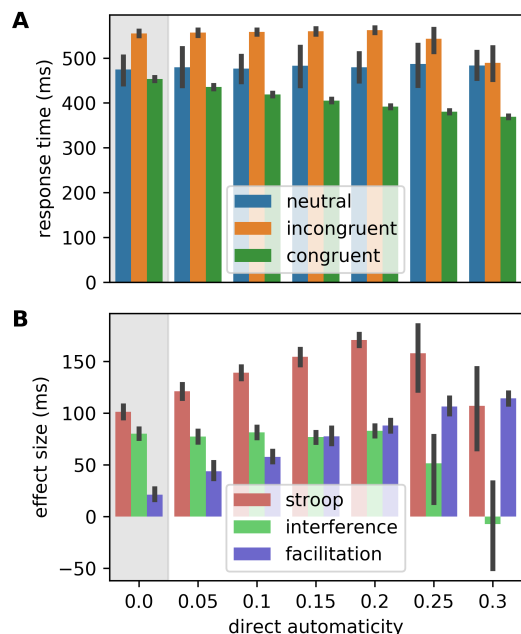


Figure 5: Model response times when introducing a direct connection between visual and result that always decodes the WORD, bypassing the attention system.

model. This includes adding a more detailed visual system, potentially modelling visual effects such as how increasing the spacing between letters in a word can reduce the Stroop effect. Furthermore, non-neural models of the Stroop effect (e.g. Lovett, 2005) have included cognitive strategies, which can be added to our model using the existing Semantic Pointer Architecture techniques.

Even with these limitations, this model presents an intriguing alternative to localist accounts of the Stroop effect. However, more work must be done to validate this as a cognitively plausible model.

References

- Ashcraft, M. H. (1994). *Human memory and cognition*. New York, NY: Harper Collins.
- Augustinova, M., Parris, B. A., & Ferrand, L. (2019). The loci of stroop interference and facilitation effects with manual and vocal responses. *Frontiers in psychology*, 10.
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological review*, 108(3), 624.
- Bugg, J. M., Jacoby, L. L., & Toth, J. P. (2008). Multiple levels of control in the stroop task. *Memory & cognition*, 36(8), 1484–1494.
- Bustamante, L. A., Lieder, F., Musslick, S., Shenhay, A., & Cohen, J. (2020). Learning to overexert cognitive control in a stroop task.
- Caron, E. E., Reynolds, M. G., Ralph, B. C. W., Carriere, J. S. A., Besner, D., & Smilek, D. (in press). Does pos-

ture influence the stroop effect? *Journal of Psychological Science*.

- Choo, X., & Eliasmith, C. (2010, 08/2010). A spiking neuron model of serial-order recall. In R. Cattrambone & S. Ohlsson (Eds.), . Portland, OR: Cognitive Science Society.
- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: a parallel distributed processing account of the stroop effect. *Psychological review*, 97(3), 332.
- Cohen, J. D., & Huston, T. A. (1994). "progress in the use of interactive models for understanding attention and performance. *Attention and performance XV: Conscious and nonconscious information processing*, 15, 453.
- Corbetta, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., & Petersen, S. E. (1991). Selective and divided attention during visual discriminations of shape, color, and speed: functional anatomy by positron emission tomography. *Journal of neuroscience*, 11(8), 2383–2402.
- Dunbar, K., & MacLeod, C. M. (1984). A horse race of a different color: Stroop interference patterns with transformed words. *Journal of Experimental Psychology: Human Perception and Performance*, 10(5), 622.
- Eliasmith, C. (2013). *How to build a brain: A neural architecture for biological cognition*. New York, NY: Oxford University Press.
- Eliasmith, C., & Anderson, C. H. (2003). *Neural engineering: Computation, representation, and dynamics in neurobiological systems* [book]. Cambridge, MA: MIT Press.
- Georgopoulos, A., Schwartz, A., & Kettner, R. (1986). Neuronal population coding of movement direction. *Science*, 233(4771), 1416–1419. doi: 10.1126/science.3749885
- Gosmann, J., Voelker, A. R., & Eliasmith, C. (2017). A spiking independent accumulator model for winner-take-all computation. In *Proceedings of the 39th annual conference of the cognitive science society*. London, UK: Cognitive Science Society.
- Herd, S. A., Banich, M. T., & O'reilly, R. C. (2006). Neural mechanisms of cognitive control: An integrative model of stroop task performance and fmri data. *Journal of cognitive neuroscience*, 18, 22–32.
- Kajić, I., Gosmann, J., Stewart, T. C., Wennekers, T., & Eliasmith, C. (2017). A spiking neuron model of word associations for the remote associates test. *Frontiers in Psychology*, 8, 99. doi: 10.3389/fpsyg.2017.00099
- Kajić, I., Schröder, T., Stewart, T. C., & Thagard, P. (2019). The semantic pointer theory of emotion: Integrating physiology, appraisal, and construction. *Cognitive Systems Research*. doi: <https://doi.org/10.1016/j.cogsys.2019.04.007>
- Levy, S. D., & Gayler, R. (2008). Vector symbolic architectures: A new building material for artificial general intelligence. In *Proceedings of the 2008 conference on artificial general intelligence 2008: Proceedings of the first agi conference* (p. 414–418). NLD: IOS Press.
- Lovett, M. C. (2005). A strategy-based interpretation of stroop. *Cognitive Science*, 29(3), 493–524. doi:

- Lu, T., Voelker, A. R., Komer, B., & Eliasmith, C. (2019). Representing spatial relations with fractional binding. In *41st annual meeting of the cognitive science society*. Montreal, QC: Cognitive Science Society.
- MacLeod, C. M. (1991). Half a century of research on the stroop effect: an integrative review. *Psychological bulletin*, 109(2), 163.
- Petersen, S. E., & Posner, M. I. (2012). The attention system of the human brain: 20 years after. *Annual review of neuroscience*, 35, 73–89.
- Plate, T. A. (1995). Holographic reduced representations. *IEEE Transactions on Neural Networks*, 6(3), 623–641.
- Rayner, K., & Pollatsek, A. (1989). *A. (1989). the psychology of reading englewood cliffs*. NJ: Prentice-Hall.
- Rumelhart, D. E., McClelland, J. L., & PDP Research Group, C. (Eds.). (1986). *Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1: Foundations*. Cambridge, MA, USA: MIT Press.
- Stewart, T. C., Bekolay, T., & Eliasmith, C. (2011). Neural representations of compositional structures: Representing and manipulating vector spaces with spiking neurons. *Connection Science*, 23(2), 145–153.
- Stewart, T. C., & Eliasmith, C. (2012). Compositionality and biologically plausible models. In W. Hinzen, M. Werning, & E. Machery (Eds.), *Oxford handbook of compositionality*. Oxford University Press.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of experimental psychology*, 18(6), 643.